

Optimal State-Space Reduction for Pedigree Hidden Markov Models

B. Kirkpatrick, K. Kirkpatrick

February 14, 2012

Abstract

To analyze whole-genome genetic data inherited in families, the likelihood is typically obtained from a Hidden Markov Model (HMM) having a state space of 2^n hidden states where n is the number of meioses or edges in the pedigree. There have been several attempts to speed up this calculation by reducing the state-space of the HMM. One of these methods has been automated in a calculation that is more efficient than the naïve HMM calculation; however, that method treats a special case and the efficiency gain is available for only those rare pedigrees containing long chains of single-child lineages. The other existing state-space reduction method treats the general case, but the existing algorithm has super-exponential running time.

We present three formulations of the state-space reduction problem, two dealing with groups and one with partitions. One of these problems, the maximum isometry group problem was discussed in detail by Browning and Browning [2]. We show that for pedigrees, all three of these problems have identical solutions. Furthermore, we are able to prove the uniqueness of the solution using the algorithm that we introduce. This algorithm leverages the insight provided by the equivalence between the partition and group formulations of the problem to quickly find the optimal state-space reduction for general pedigrees.

We propose a new likelihood calculation which is a two-stage process: find the optimal state-space, then run the HMM forward-backward algorithm on the optimal state-space. In comparison with the one-stage HMM calculation, this new method more quickly calculates the exact pedigree likelihood.

1 Introduction

Motivation. Statistical calculations on pedigrees are the principal method behind the most accurate disease-association approaches [15, 18]. In those approaches, the aim is to find the regions of the genome that are associated with the presence or absence of a disease among related individuals. Furthermore, pedigree likelihoods are used to estimate fine-scale recombination rates in humans [4], where there are few other approaches for making these estimates. Many implementations of the likelihood estimates for pedigrees exist [7, 1, 16]. Estimates of probabilities on pedigrees are of great interest to computer scientists because they give an important example of graphical models which model probability distributions by using a graph whose edges are conditional probability events and whose nodes are random variables [12]. Methods for reducing the state-space of a pedigree graphical model could generalize to other graphical models, as noted also by Geiger et al [8].

The Problem Summary. Hidden Markov Models (HMMs) analyzing the genotypes of related individuals have running time $O(m2^n)$ where m is the number of sites and n is the number of meioses in the pedigree. Therefore, it is desirable to find more efficient algorithms. Any partitioning of the state space into k ensemble states (i.e., states with identical emission probabilities and Markovian transition probabilities) will improve the running time of an HMM to $O(mk)$, even if the ensembles are not optimal. Since the HMMs have an exponential state space and a running time polynomial in the size of the state space, even an exponential algorithm for finding ensemble states can improve the running time of the HMM calculations.

Literature Review. Donnelly [5] introduced the idea of finding ensemble states for the IBD Markov model, and used a manual method for finding the symmetries for several examples of two-person pedigrees. Browning and Browning [2] formalized the requirements for symmetries that describe ensemble states in a new HMM. They gave the first algorithm for finding the maximal set of isometries that preserves the Markov property and the IBD information. Their algorithm which is based on enumerating permutations appears to have worst-case running time of at least of $O(n!2^{2n})$, where n is the number of meioses in the pedigree. However, the running time of their algorithm is difficult to analyze. They also left open the question of whether groups other than isometry groups could give useful state-space reductions [2].

McPeck [14] presented a detailed formulation of the condensed identity states and an algorithm. Most recently Geiger et al [8] discussed a similar problem using the language of partitions. In this paper, we will prove that the partition problem and the isometry problem are identical. Geiger, et al. gave a special-case state-space reduction involving only partitions that collapse simple lineages (multiple generations with a single child per generation and with the non-lineage parents being founders). Several people have introduced algorithms for finding symmetries for systems applications [13, 10].

Our Contribution. Inspired by the work of Browning and Browning [2], we look for maximal ensembles of the hidden states that can be used to create a new HMM with a much more efficient running-time. We introduce an improved algorithm for finding the maximal ensemble states, sets of hidden states, that preserve both the Markov property and the identity by descent (IBD) information of the individuals of interest.

Related work includes the algorithms of Geiger, et al [8], and Browning and Browning [2]. Geiger, et al. found isometries of a limited type in $O(n^2)$ time and demonstrated their practical value. Browning and Browning introduced a super-exponential algorithm that finds a maximal group of isometries, i.e. the largest number of elements. However, they did not draw any conclusions about whether their method finds the group with the maximal orbit sizes.

We introduce an $O(n2^{2n})$ maximal-ensemble algorithm for finding a permutation group on the 2^n vertices of the hypercube, and for producing the most efficient ensemble states (i.e. the smallest partition of the state-space that respects the IBD and Markov properties and has the maximal partition sets and minimal number of sets in the partition). We prove that the optimal partition is a solution to the maximal isometry group problem that Browning and Browning introduced. Therefore, both their algorithm and ours finds the optimal partition of the state space which can be described using a group of isometries having a maximal number of elements. However, our algorithm is much faster, having a coefficient n instead of $n!$.

We also introduce a bootstrap version of the maximal-ensemble algorithm which takes advan-

tage of the isometries introduced by Geiger, et al. [8] and the well-known founder isometry. By enumerating one representative from each set of the partition induced by the known isometries, we can create a bootstrap maximal-ensemble algorithm that runs in $O(nk2^n)$ time where n is the number of meioses in the pedigree, and k is the number of partitions from the known isometries.

2 Problem Description

Consider a pedigree graph, P , having individuals V as nodes and having n meioses with each meiosis being a directed edge from parent to child. Let I being the set of individuals of interest, perhaps because we have data for those individuals. While it might be algorithmically convenient to assume that $I = V$, it is impractical. Many of the ancestral individuals in the pedigree are likely deceased, and genetic samples are unavailable.

An *inheritance state or vector* is a binary vector x with n bits where each bit indicates which grand-parental allele, paternal or maternal, was copied for that meiosis. The equivalent *inheritance graph*, R_x , has two nodes per individual (one for each allele) and edges from inherited parental alleles to their corresponding child alleles. Individuals of interest are called *identical by descent (IBD)* if a particular founder allele was copied to each of the individuals. In general, the inheritance graph is a collection of trees, since each allele is copied from a single parent.

The set of all inheritance states (binary n -vectors) is the n -dimensional hypercube \mathcal{H}_n , with 2^n vertices. The inheritance process is modeled as a symmetric random walk on \mathcal{H}_n , with the time dimension of the walk being the distance along the genome. At equilibrium, the walk has uniform probability of being at any of the hypercube vertices. From vertex x in \mathcal{H}_n , a step is taken to a neighboring vertex after an exponential waiting time with parameter $\lambda = n$. For each individual zygote, with one meiosis, this is a Poisson process with parameter $\lambda = 1$ and genome length roughly 30.

There is a discrete version of this random walk, which is often used for hidden Markov models (HMMs) that compute the probability of observing the given data by taking an expectation over the possible random walks on the hypercube. Let X be a Markov process, $\{X_t : t = 1, 2, \dots, m\}$ for m loci with a state space \mathcal{H}_n consisting of all the inheritance states of the pedigree. The recombination rate, θ_t , is the probability of recombination per meiosis, between a neighboring pair of loci, t and $t + 1$. If t and $t + 1$ are separated by distance d , then the Poisson process tells us that the probability of an odd number of recombinations is $\theta_t = 1/2(1 - e^{-2\lambda d})$. The natural distance on \mathcal{H}_n is the Hamming distance, $|x \oplus y|$, for two states x and y , where \oplus is the XOR operation and $|\cdot|$ is the L^1 -norm in \mathbb{R}^n . Then the probability of transitioning from x to y is

$$Pr[X_{t+1} = y | X_t = x] = \theta_t^{|x \oplus y|} (1 - \theta_t)^{n - |x \oplus y|}.$$

Figure 2 shows an example HMM with three genomic sites. The states of the HMM are shown in circles on the right.

We define potential ensembles of states as being the orbits of a group of symmetries. Let G be a group that acts on the state space \mathcal{H}_n of X . A symmetry is a bijection $\psi \in G$ where ψ is a permutation on 2^n elements, the vertices of \mathcal{H}_n . An *orbit* of G acting on \mathcal{H}_n is the set

$$\omega(y) = \{x | x = \psi(y) \text{ and } \psi \in G\},$$

and we write the set of all orbits of G as $\Omega(G) = \{\omega(y) : y \in \mathcal{H}_n\}$.




Figure 1: **Two Half-Siblings.** (Left Panel) A pedigree with two non-founders of which two are half-siblings together with their common parent. Circles and boxes represent female and male individuals, while the two black dots for each person represent their two chromosomes or alleles. Edges are implicitly directed downward from parent to child. The alleles of each individual are ordered, so that the left allele, or paternal allele, is inherited from the person's father, while the right, maternal allele is inherited from the mother. The two siblings are the only labeled individuals. Their genomes are shown in color so that the same color indicates inheritance from the same ancestor. For convenience, the genotype of each person is homozygous. (Right Panel) The HMM for the genotypes from the left panel. At each site in the genome, the possible states are the vectors in \mathcal{H}_n . In each circle an inheritance state is drawn as an inheritance graph and the inheritance states for a single site are arranged in a column. The allowed transitions between neighboring sites are a complete bipartite graph (due to space, only a fraction of the edges are drawn). The nodes with a slash through them are inheritance states that are not allowed by the data. The red nodes and edges are the path for the actual inheritance states indicated by the yellow and blue in the left panel. However, this is only one of several paths of inheritance states that are consistent with the data.

Conventional algorithms for computing likelihoods of data have an exponential running time, because the state space of the HMM is exponential in the number of meioses in the pedigree. We propose new ways to collapse hypercube vertices into ensemble states for a new HMM that has a more efficient running time. In particular we are interested in optimal ensemble states that preserve certain relationship structures: the Markovianity of the random walk and the emission probabilities. We will first discuss the Markov property and then discuss the constraints on ensemble states that the emission probabilities provide.

2.1 Markov Property

Let $\{X_t\}$ be a stationary, reversible Markov chain with state space Ω , such as the chain corresponding to the hidden states of the pedigree HMM.

Let Y be a new process, $\{Y_t : t = 1, 2, \dots, m\}$ having states $\Omega(G) = \{\omega_1, \dots, \omega_k\}$ which are the orbits of some group G . This new Markov chain is coupled to the original such that when

$X_t = x \in \omega \in \Omega(G)$, $Y_t = \omega$, and Y_t is a projection of X_t into a smaller state space. Define the transition probabilities for process Y_t as

$$Pr[X_{t+1} \in \omega_j | X_t = x] = \sum_{y \in \omega_j} Pr[X_{t+1} = y | X_t = x] \quad \text{for } x \in \omega_i, \text{ for } \omega_i, \omega_j \in \Omega(G). \quad (1)$$

We will call Y_t the expectation chain since

$$Pr[Y_{t+1} = \omega_j | Y_t = \omega_i] = \mathbb{E}[E_j | X_t = x],$$

where E_j is the event that $X_{t+1} \in \omega_j$.

Since X_t is stationary and reversible, the necessary and sufficient condition [3] for Y_t to also be Markov is that

$$\sum_{y \in \omega_j} Pr[X_{t+1} = y | X_t = x_1] = \sum_{y \in \omega_j} Pr[X_{t+1} = y | X_t = x_2] \quad (2)$$

for all $x_1, x_2 \in \omega_i$ for all i , and for all ω_j . Therefore any group whose orbits satisfy this set of equations can be used to create a new Markov chain Y_t .

From Equations (1) and (2) we see that the stationary distribution of Markov chain Y_t is $Pr[Y_t = \omega_i] = \sum_{y \in \omega_i} \pi_y$ where π_y is the stationary distribution of X_t . For pedigree HMMs, the stationary distribution of X_t is uniform, $\pi_y = 1/2^n$, therefore the expectation chain for some group that satisfies Equation (2) will have a stationary distribution $Pr[Y_t = \omega_i] = |\omega_i|/2^n$.

For pedigree Markov chains, Equation (2) becomes, for $s = \theta/(1 - \theta)$ and $0 < \theta < 0.5$,

$$\sum_{y \in \omega_j} s^{|y \oplus x_1|} = \sum_{y \in \omega_j} s^{|y \oplus x_2|} \quad \forall x_1, x_2 \in \omega_i. \quad (3)$$

If the expectation chain Y_t corresponding to pedigree Markov chain X_t satisfies this equation, we say that it satisfies the *Markov property*. Notice that these polynomials are identical if and only if the coefficients of like powers are equal.

Browning and Browning [2] gave an algorithm that searches for a maximal group of isometries where the group was maximal in the number of group elements. A group, G , of isometries has orbits $\Omega(G) = \{\omega_1, \dots, \omega_k\}$ such that $|T(x) \oplus T(y)| = |x \oplus y|$ for all $T \in G$, $y \in \omega_j$ and $x \in \omega_i$ for all i, j . This means that the transition probabilities are related by

$$Pr[X_{t+1} = y | X_t = x] = Pr[X_{t+1} = T(y) | X_t = T(x)]. \quad (4)$$

Browning and Browning left open the question of whether any symmetry groups satisfying Equation (3) must be equivalent to a group of isometries (meaning that it has the same orbits). We answer this question. Theorem 1 proves that for any permutation satisfying Equation 3, there is always a group of isometries having the same orbits as the permutation sets.

Theorem 1. *Let S be a group of permutations of \mathcal{H}_n whose orbits $\Omega(S)$ satisfy Equation (3). Then there exists a group of isometries G having the same orbits as S : that is, for every $T \in G$ and all $x, y \in \mathcal{H}_n$, $|y \oplus x| = |T(y) \oplus T(x)|$, and the set of orbits $\Omega(G)$ is equal to $\Omega(S)$.*

Proof. We prove this by constructing a generating set A for G . First, let the identity permutation π_e be in A . Then for each orbit ω of S , and each pair of points x_1 and x_2 in ω , we will construct a permutation π_{x_1, x_2} to add to the generating set A . If $x_1 = x_2$, then $\pi_{x_1, x_2} = \pi_e$ which is already

in A . If $x_1 \neq x_2$ then π_{x_1, x_2} will be a composition of disjoint two-cycles, in particular including the cycle $(x_1 \ x_2)$. Our generating set A will then be the union of all these permutations, so by construction it will generate a group $G = \langle A \rangle$ having the same orbits as S .

For fixed $x_1, x_2 \in \omega$, the two-cycles comprising π_{x_1, x_2} are constructed as follows:

For each $k = 1, \dots, n$, define $a_k := \#\{y \in \omega : |y \oplus x_1| = k\}$ and $b_k := \#\{z \in \omega : |z \oplus x_2| = k\}$, which implies by Equation (3) that $a_k s^k = b_k s^k$ for each k , and hence $a_k = b_k$, since $s > 0$ and polynomials in s are uniquely determined by their coefficients and powers. Then, for each given $y_1 \in \omega$ that is distinct from both x_1 and x_2 , there exists z_1 such that $|y_1 \oplus x_1| = |z_1 \oplus x_2| = k$, because $a_k \geq 1$, a consequence of the fact that $y_1 \in A_k := \{y \in \omega : |y \oplus x_1| = k\}$. In other words, $z_1 := y_1 \oplus (x_1 \oplus x_2)$, and the cycle is $c_1 := (y_1 \ z_1)$.

Proceed similarly for $y_2 \in \mathcal{H}_n \setminus \{x_1, x_2, y_1, z_1\}$, defining $z_2 := y_2 \oplus (x_1 \oplus x_2)$, and the cycle $c_2 := (y_2 \ z_2)$, and so on for each $y_i \in \mathcal{H}_n \setminus \{x_1, x_2, y_1, z_1, \dots, y_{i-1}, z_{i-1}\}$, with $z_i := y_i \oplus (x_1 \oplus x_2)$ and $c_i := (y_i \ z_i)$. Then we define the permutation $\pi_{x_1, x_2} := c_1 \circ c_2 \circ \dots \circ c_{2^n}$. In particular it has the cycle $(x_1 \ x_2)$ in its composition, since when $y = x_1$, we have $z = x_2$. Notice also that the definitions of z_i imply that

$$y_i \oplus y_j = y_i \oplus y_1 \oplus y_j \oplus y_1 = z_i \oplus z_1 \oplus z_j \oplus z_1 = z_i \oplus z_j; \quad (5)$$

$$y_i \oplus z_j = y_i \oplus x_1 \oplus z_j \oplus x_1 = z_i \oplus x_2 \oplus y_j \oplus x_2 = z_i \oplus y_j. \quad (6)$$

Hence by taking L^1 norms, the permutation π_{x_1, x_2} is an isometry with respect to Hamming distance.

Furthermore, the group $G = \langle A \rangle$ will have the same orbits as S , since for each orbit ω and each pair $x_1, x_2 \in \omega$, the cycle $(x_1 \ x_2)$ will appear in some permutation, and no pair of points from different orbits will appear as a cycle in any permutation. \square

This proof complements the result from Browning and Browning regarding the fact that isometry groups always satisfy Equation 3. Indeed, we will state the complete result as a corollary.

Corollary 2. *A group S has orbits $\Omega(S)$ satisfying Equation 3 if and only if there is an isometry group G whose orbits $\Omega(G)$ are identical to $\Omega(S)$.*

Proof. Browning and Browning [2] showed that all isometry groups G satisfy Equation 3. Theorem 1 completes the proof. \square

It is a well-known fact in algebra that any partition can be the orbits of some symmetry group, and that the orbits of any symmetry group are a partition [6]. We will recapitulate this simple result next.

Corollary 3. *A partition satisfies Equation 3 if and only if it is equivalent to the orbits of some isometry group.*

Proof. Assume we are given a partition $\{W_1, \dots, W_k\}$ of set \mathcal{H}_n where $W_i \cup W_j = \emptyset$, $\cup_i W_i = \mathcal{H}_n$ and the partition satisfies Equation 3. We will create a symmetry group S whose orbits $\Omega(S) = \{W_1, \dots, W_k\}$. This is easily done. For each set in the partition W_i , create a permutation with a single cycle $\pi_i = (y_1 \ y_2 \ \dots \ y_l)$ where all $y_j \in W_i$. Make a generating set $A = \{\pi_i : 1 \leq i \leq k\} \cup \pi_e$ where π_e is the identity permutation. Then group $S = \langle A \rangle$ clearly has orbits $\Omega(s) = \{W_1, \dots, W_k\}$. By Theorem 1 there is an isometry group with the same orbits.

Assume we are given an isometry group G . Clearly, by Browning and Browning's proof [2], the orbits define a partition $\Omega(G)$ that satisfies Equation 3. \square

Browning and Browning [2] also showed that any isometry $T : \mathcal{H}_n \rightarrow \mathcal{H}_n$ can be uniquely written as $T = \pi \circ \phi_a$ where π is a permutation on n elements, the bits of the hypercube vertex, and ϕ_a is a switch function where $\phi_a(x) = a \oplus x$ where \oplus is the bit-wise XOR operation.

An isometry describes some aspect of the pedigree graph. For example, an isometry consisting of a switch and the identity permutation can be used to enumerate one element from each orbit by simply fixing the 1-bit's value and then enumerating all possible values for the other switch bits. On the other hand, an isometry consisting of the identity switch (all zero) and a permutation of one cycle can be used to enumerate one element for each orbit by listing the 1-prefixes of the permuted bits, (i.e. for three bits, the representatives are 000, 100, 110, and 111).

2.2 Emission Property

The Markov property is not enough to ensure that the HMM based on Markov chain Y_t has the same likelihood as the HMM for X_t . In order to ensure this, we introduce a property on the emission probabilities, namely that all the elements in one orbit must have identical emission probabilities. We call these orbits the emission partition, since they are induced by the emission probability. In order to define this object, we need to introduce some more notation.

Recall that R_x is the inheritance graph for inheritance path x . The relationship structures we wish to preserve are the IBD relationships on the individuals of interest I . Let I_m be the maternal alleles of all the individuals of interest and I_p be the paternal alleles of all the individuals of interest. The inheritance graph R_x is a forest of trees; let $CC(R_x)$ refer to the connected components of R_x which are labeled with $I_m \cup I_p$. The same-labeled connected components induce a partition

$$D_x = \{y \in \mathcal{H}_n | CC(R_y) = CC(R_x)\}.$$

We call the partition $\{D_x | \forall x\}$ the identity states, since it indicates a particular identity-by-descent (IBD) relationship among the labeled individuals. These have been well studied [9, 17, 11].

Looking at a small example, containing two siblings who are the individuals of interest and their two parents, we see that the identity states are:

$$D_{0000} = \{0000, 0101, 1010, 1111\},$$

$$D_{1000} = \{1000, 0010, 1101, 0111\},$$

$$D_{0100} = \{0100, 1110, 0001, 1011\},$$

$$D_{1100} = \{1100, 0110, 1001, 0011\}.$$

But if we think carefully about this example, there is symmetry in the pedigree, namely swapping the two parents, that does not appear in this partition. Due to this reason, we need to consider the following object.

Let $Pr[D|X_t]$ be the probability that the state X_t of the HMM emits the observed data D in state t . Then the partition induced on the state space by the emission probability is the *emission partition* containing all distinct sets E_x :

$$E_x = \{y \in \mathcal{H}_n \mid Pr[D = d | X_t = x] = Pr[D = d | X_t = y] \ \forall d\}$$

and

$$Pr[D = d | X_t = x] = \sum_{\tilde{d} \text{ consistent with } R_y} \frac{1}{2^{h(d)}} \prod_{c \in CC(R_x)} Pr[c(\tilde{d})]$$

where d is a vector of sets, \tilde{d} is a vector of tuples that is an ordered version of d , meaning that $d_i \equiv \tilde{d}_i$ while removing the order information from \tilde{d}_i , and $c(\tilde{d})$ gives the allele of \tilde{d} that is assigned to that connected component, and $h(d)$ is the number of heterozygous sites in d . Note that each connected component is a tree, and has exactly one founder. Also, the identity states are consistent with these probabilities, but the identity states are a sub-partition of the emission partition. Specifically, from our previous example, $0100 \notin D_{1000}$, but $0100 \in E_{1000}$. Indeed, the emission partition for the example is $D_{0000}, \{D_{1000}, D_{0100}\}, D_{1100}$.

We say that the expectation Markov chain Y_t satisfies the *emission property* if and only if it preserves the emission partition in order for the corresponding HMM to have the correct likelihood. To preserve the emission partition, all the group elements $T \in G$ must satisfy $T(y) \in E_x$ for all $y \in E_x$ and for all x .

Now, it is necessary to compute the E_x quickly. The naïve algorithm would be slow, since we would have to consider all pairs x, y and all possible data d . Neither can we use the methods in the literature dealing with *condensed identity states* [9, 17, 11], because the literature computes pedigree-free condensed identity states. That calculation takes the sets from the identity states and applies permutations of the form $\pi_i = (i_m \ i_f)$ to swap the alleles of an individual of interest $i \in I$. However these permutations can violate the inheritance rules specified by a fixed pedigree. For the example above, take vector $1010 \in D_{0000}$ and swap the alleles of the second child $\pi_2(1010) = 1001 \in D_{1100}$. This clearly produces a partition that is not the emission partition, and so it would violate the property that we wish to enforce. Several works on optimal state space reduction for pedigree HMMs have discussed the condensed identity states [2, 14] for state-space reduction. It would appear that they did not formulate the emission partition that was mentioned by Geiger, et al. [8] and that is used here.

The main difference between D and E partitions is that the probability $Pr[D = d | X_t = x]$ has a product over indistinguishable connected components, whereas the identity states distinguishes each connected component. This partition must answer the question of which connected components are exchangeable. Let I' be the individuals of interest having parents who are not individuals of interest. So, we can rewrite E_x as follows:

$$E_x = \{y \in \mathcal{H}_n \mid \exists \phi \text{ a proper isomorphism s.t. } CC(R_x) = CC(\phi(R_y))\}$$

where an isomorphism ϕ is *proper* if and only if ϕ is an isomorphism from R_y to R_x where for all $i \in I' \cup V \setminus I$, either $\phi(i_f) = i_f$ and $\phi(i_m) = i_m$ or $\phi(i_f) = i_m$ and $\phi(i_m) = i_f$. This definition of E_x is easier to compute, because now we can do an $O(n)$ check to see if the forest of trees in x and y are isomorphic, which leads to an $O(n2^{2n})$ calculation. However, we can do better.

From the above definition, we see that in order for two inheritance paths to be isomorphic, the pedigree graph itself (as opposed to the inheritance graph) must have a automorphism. If we can identify all the relevant automorphisms for the pedigree graph, then we can make a set A of permutations (one for each automorphism), and use a group theoretic algorithm for obtaining the orbits of $\langle A \rangle$ acting on the partition $\{D_x \mid \forall x \in \mathcal{H}_n\}$ to obtain the desired emission partition.

First to obtain the automorphisms of the graph, we will employ a naïve strategy. Let $i \in I' \cup V \setminus I$ be an individual of interest. Recall that any proper isomorphism must map one branch

of i 's ancestral lineage to the other branch. In order to be consistent, for the set $J = \{i\} \cup \{j \mid j \text{ full sib of } i\}$, the automorphism must $\phi(j_m) = j_f$ for $j \in J$. Considering i 's parents and proceeding backward in time, the sub-pedigree connected to the ancestors forms a directed acyclic graph (dag) with in-degree two. Without loss of generality, we can assume that this sub-pedigree has no individuals in $I \setminus \{i\}$, because if there were there would be no proper automorphism, and there is a descendant of the ancestors not in I it can be trivially removed from the pedigree [14]. Therefore, we may consider only the tree of direct ancestors branching backward in time. At each branch point, b , in this tree, we assign an indicator $\gamma_b = 1$ if the father is to the left and the mother to the right. There are $O(2^n)$ assignments of these variables $\{\gamma_b \mid \forall b\}$. For each possible assignment, perform an $O(n)$ graph-traversal operation to check whether this assignment is an automorphism. We take the first automorphism ϕ that we find, because any other ϕ' from the same lineage will satisfy $CC(\phi(R_x)) = CC(\phi'(R_x))$ for all inheritance paths x .

Now that we have the automorphisms, we can write them as isometries and put them in set A and consider the orbits of the group $\langle A \rangle$ acting on the identity states. These orbits are the emission partition. To obtain these orbits, we will use the well-known orbit algorithm from computational group theory. Notice, that we wish to apply this algorithm to the existing partition $M := \{D_x \mid \forall x\}$. Take one set $D_x \in M$ and initialize its orbit as $O_x := \{D_x\}$. At the end of the following procedure O_x will contain all the elements in x 's orbit. For every element $D_x \in O$ and every automorphism permutation $a \in A$, compute $y := a(z) \quad \forall z \in D_x$. If $y \notin O_x$, then this y and all the elements in its set D_y are added to O_x and D_y is removed from M . This procedure is repeated until M is empty. Notice that $CC(y) = CC(a(z))$ is compared to $CC(x)$ to determine if y is in O_x .

Since the comparison $CC(y) = CC(x)$ can be computed in linear time. The running time to obtain the automorphisms is $O(n2^n)$ and the orbit algorithm runs in $O(n2^n)$ time. This means that obtaining automorphisms of the pedigree is preferable to checking pairs of inheritance vectors for isomorphism.

2.3 Examples

We will consider two examples, here. The first is a three generation pedigree while the second is a result that applies to all two generation pedigrees.

2.3.1 Three-Generation Pedigree

For example, given 4 meioses for two half-cousins, A and B , with one shared grandparent, their common grandparent and their respective parents who are half-siblings, we have 16 hypercube vertices (see Figure 2.3.1). Our individuals of interest are $I = \{A, B\}$. The emission partition is, in this case, identical to the identity states and contains the sets $E_1 = \{A_p\}\{A_m, B_m\}\{B_p\}$ and $E_2 = \{A_p\}\{A_m\}\{B_m\}\{B_p\}$, since these are the only partitions of alleles of individuals I that have non-empty sets in the emission partition. The emission partition induced on the hypercube vertices is: $E_{x_1} = \{1001, 1111\}$ and $E_{x_2} = \mathcal{H}_n \setminus E_{x_1}$.

Notice that in this instance we cannot use the emission partition $\{E_x \mid \forall x\}$ as the state space of a new Markov chain. For example, if we were to let Z_t be a Markov chain on the partition given by the emission partition, then the Markov criteria would fail to hold. Specifically, consider state $x_1 = 0001$ and $x_2 = 0011$. Then by checking Equation (2), we have $\sum_{y \in E_{x_1}} Pr[X_t = y \mid X_t = 0001] = \theta(1 - \theta)^3 + \theta^3(1 - \theta)$ but $\sum_{y \in E_{x_1}} Pr[X_t = y \mid X_t = 0011] = 2 \cdot \theta^2(1 - \theta)^2$.

The largest partition of \mathcal{H}_n that satisfies the Markov criteria is $P_J = \{1001, 1111\}$, $P_R = \{0010, 0100\}$, $P_G = \{1011, 1101\}$, $P_B = \{0000, 0110\}$, $P_K = \{0011, 0101, 1010, 1100\}$, and $P_L = \{0001, 0111, 1000, 1110\}$. Let H be the matrix of pair-wise Hamming distances between all the vertices of the hypercube. Then the transition probabilities take the form: For example, $Pr[Y_{t+1} = P_L | Y_t = P_K] = 2\theta(1 - \theta)^3 + 2\theta^3(1 - \theta)$.




Figure 2: **Two Half-Cousins.** (Left Panel) A pedigree with four non-founders of which two are half-cousins together with their common grandparent. As before, the two black dots for each person represent their two alleles, and the alleles of each individual are ordered, so that the left allele, or paternal allele, is inherited from the person's father, while the right, maternal allele is inherited from the mother. The two cousins are labeled A and B . It is easy to see that the only possible IBD is between alleles A_m and B_m , the maternal alleles of individuals A and B , respectively. (Right Panel) This makes the four male founders irrelevant to the question of IBD. The four meioses are labeled in the order of their bits, left-to-right, and the inheritance states are represented in binary as $x_1x_2x_3x_4$. Let $x_i = 0$ if that allele was inherited from the parent's paternal allele, and $x_i = 1$ if from the maternal allele. For instance, A and B are IBD only for inheritance states 1001 and 1111.

Notice that this partition can be expressed as the orbits of a group of isometries, because $G = \langle (1\ 4), (2\ 3), \phi_{0110} \rangle$ does not violate the IBD class.

2.3.2 Two-Generation Pedigrees

Lemma 4. *For any two-generation pedigree, the partition defined by the emission partition, $C = \{E_x | \forall x\}$, satisfies the Markov Property.*

Proof. We can establish this by finding a group of isometries whose orbits are the emission partition. This group has the generating set A where $A = \{\phi_f : \forall f\} \cap \{\pi_m : \forall m\}$ and ϕ_f and π_m are defined as follows. For founder f , ϕ_f is a switch having bits set as follows. Let i_1, \dots, i_c be the meioses from founder f to each of the founders c children. Then $\phi_{fi} = 1$ if $i = i_j$ for some j and $\phi_{fi} = 0$ otherwise. Let $m = (f_1, f_2)$ which are untyped monogamous married founding pairs. Then $\pi_m = c_1 \circ c_2 \circ \dots \circ c_k$ is a permutation composed of k disjoint cycles, one for each child. For child i with meiosis bits i_0, i_1 , $c_i = (i_0\ i_1)$. The group of isometries $G = \langle A \rangle$.

Now, we simply need to establish that the emission partition C are the orbits of this group G . There is no element $T \in G$ that maps $x \in E_{x_1}$ to $y \in E_{x_2}$, since every ϕ_f and π_m map the bits of x in ways that maintain $CC(R_x)$. Now, we simply need to show that for any $x_1, x_2 \in E_x$, there is always some element $T \in G$ such that $y = T(x)$. Consider each connected component in $CC(R_x)$

where x and y differ. The alleles connected in this connected component must all share inheritance through one of the founder bits of the common parents. If there is only one common parent, the switch for that founder must map between x and y in the bits for that connected component. If there are two common parents, then there must exist a composition of two founder switches and the founder permutation that maps between x and y for the bits in that connected component. The complete map T is simply the composition of the isometries for each connected component. \square

In the next section we will introduce the Maximal Ensemble Problem, and we will soon see that this lemma provides a fast method to obtain the optimal partition for two-generation pedigrees.

2.4 The State-Space Reduction Problem

There have been three state-space reduction problems posed, we restate these here. Given the original pedigree state space \mathcal{H}_n , there are three ways to reduce the state space.

Maximum Ensemble Problem [8] Find the partition, $\{W_1, \dots, W_k\}$ of \mathcal{H}_n that satisfies both the Markov property and the emission property and that maximizes the sizes of the sets:

$$\max \sum_{i=1}^k |W_i|.$$

Maximum Isometry Group Problem [2] Find the isometry group G of maximal size whose orbits $\Omega(G)$ satisfy the emission property.

Maximum Symmetry Group Problem Find the symmetry group G of maximal size whose orbits $\Omega(G)$ satisfy both the Markov property and the emission property.

We have already proven that all symmetry groups that satisfy the Markov property have an isometry group with equivalent orbits. This means that the later two problems are identical. Indeed since these last two problems are equivalent, we will refer to them collectively as the **Maximum Group Problem**. The remaining question is the relationship between the maximum ensemble problem and the maximum isometry group problem. We will first introduce a Maximum Ensemble Algorithm and use it to prove that the solution to the Maximum Ensemble Problem is unique. Using the uniqueness result, we will be able to prove the equivalence of the Maximum Ensemble and Maximum Isometry Group Problems.

3 Maximum Ensemble Algorithm

We will introduce an algorithm that solves the Maximum Ensemble Problem and that helps us establish the uniqueness of the solution. We will then use uniqueness to prove that this algorithm also solves the Maximum Group Problem.

Consider the emission partition containing, E_x for all x of interest. Of course the sets in the emission partition are disjoint. Consider the $(2^n)!$ permutations on the vertices of the hypercube. Naively, these are all candidate permutations for our group, if we wish to find the maximal group. However in this section, we focus on finding the partition that yields the maximum ensemble solution. Given the state space, the partition can be found in linear time.

We do this by iteratively sub-partitioning the partition according to the coefficients and powers appearing in Equation 3. See Algorithm 1: Bipartition which takes as input a subpartition of the emission partition. This recursion is possible since the Markov property must produce a partition

that is a sub-partition of the emission partition (i.e. in order to respect the emission partition). Indeed, as shown in Lemma 5, any pair of vectors x_1, x_2 that violate the Markov property must appear in separate sets of the partition. This recursive approach will at worst produce a partition with each element in its own set.

Algorithm 1: Bipartition only needs to compute the $2^n \times 2^n$ matrix of distances between IBD vectors, as well as do some bookkeeping. So, the total running time is $O(2^{2n})$. Since the iterative sub-partitioning at minimum splits sets in two and does not introduce new inequalities, the number of iterations of the partition algorithm is $O(\log(2^n)) = O(n)$. One iteration of Algorithm 1: Bipartition requires $O(2^{2n})$ time for each iteration, since we have to check the $2^n \times 2^n$ matrix of distances between partition elements. So, the total running time is $O(n2^{2n})$.

Now, we need to establish the correctness and uniqueness of the partition.

Lemma 5. *Let W_i, W_j be two sets of the partition such that $x_1, x_2 \in W_i$ and x_1, x_2 violate the Markov property in Equation 3, i.e. such that*

$$\sum_{y \in W_j} s^{|y \oplus x_1|} \neq \sum_{y \in W_j} s^{|y \oplus x_2|}.$$

Then even if W_j is subdivided, x_1, x_2 continue to violate Equation 3.

Proof. This is proven by a simple property of polynomials. Since

$$\sum_{y \in W_j} s^{|y \oplus x_1|} \neq \sum_{y \in W_j} s^{|y \oplus x_2|},$$

there must be at least one power for which the polynomial coefficients disagree. Let a_k and b_k be the coefficients from the left- and right-hand sides respectively. Let $A(k) = \{y : |y \oplus x_1| = k\}$, so that $a_k = |A(k)|$, and let $B(k) = \{y : |y \oplus x_2| = k\}$, so that $b_k = |B(k)|$. Let C, D be any bipartition of W_j . Therefore C and D induce a partition of $A(k)$ and $B(k)$. Specifically $A(k)$ is partitioned into sets $A(k) \cap C$ and $A(k) \cap D$, while $B(k)$ is partitioned into $B(k) \cap C$ and $B(k) \cap D$. Since $|A(k)| \neq |B(k)|$, then at least one of

$$|A(k) \cap C| \neq |B(k) \cap C|$$

or

$$|A(k) \cap D| \neq |B(k) \cap D|.$$

Therefore at least one of

$$\sum_{y \in C} s^{|y \oplus x_1|} \neq \sum_{y \in C} s^{|y \oplus x_2|},$$

or

$$\sum_{y \in D} s^{|y \oplus x_1|} \neq \sum_{y \in D} s^{|y \oplus x_2|}.$$

□

Lemma 6. *(Loop Invariant.) Once C_{i0} is added to P' , it is never subdivided again in any iteration. This is equivalent to stating the invariant that for any i ,*

$$\sum_{y \in W_j} s^{|y \oplus x_1|} = \sum_{y \in W_j} s^{|y \oplus x_2|} \quad \forall x_1, x_2 \in C_{i0} \quad \forall W_j \in P'$$

Algorithm 1 Bipartition(P) in $O(2^{2n})$ time

input:
 P : current subpartition of the emission partition

output:
 P' : violates fewer equations of the Markov property

main:
 $P' = \emptyset$
foreach $W_i \in P$ **do**
 $C_{i0} = W_i$
 $C_{i1} = \emptyset$
foreach $W_j \in P$ **do**
 $a_k = 0$ for all $0 \leq k \leq n$
 $s_{x'} = 0$ for all $x' \in C_{i0}$

Let $x_1 \in C_{i0}$ be a fixed element of C_{i0} .

foreach $x \in C_{i0}$ **do**
 $b_k = 0$ for all $0 \leq k \leq n$
foreach $y \in W_j$ **do**

Let $k = |y \oplus x|$
if $x == x_1$ **then**
 $a_k ++$
end if
 $b_k ++$
end for
if $a_k \neq b_k$ for some $0 \leq k \leq n$ **then**
 $s_x = 1$
end if
end for

{Bipartition W_i }

foreach $x \in C_{i0}$ **do**
 $C_{i0} \leftarrow C_{i0} \setminus \{x\}$
 $C_{s_x} \leftarrow C_{s_x} \cup \{x\}$
end for
end for
 $P' \leftarrow P' \cup \{C_{i0}, C_{i1}\}$
end for

RETURN P'

Proof. Notice that the above invariant is a consequence of both the loop “foreach $W_j \in P$ ” and of the Bipartition algorithm. For the base case $C_{i0} = \emptyset \ \forall i$, and the invariant holds trivially.

Now we need to inductively prove that the invariant holds. Assume that for some i , the invariant holds. Now, consider the loop for a fixed $W_j \in P$. W_j may be partitioned into some C_{j0} and C_{j1} . Our task is to prove that for the new partition of W_j the invariant holds, i.e. that

$$\sum_{y \in C_{j0}} s^{|y \oplus x_1|} = \sum_{y \in C_{j0}} s^{|y \oplus x_2|} \quad \forall x_1, x_2 \in C_{i0}.$$

From the invariant, we have $\sum_{y \in W_j} s^{|y \oplus x_1|} = \sum_{y \in W_j} s^{|y \oplus x_2|} \quad \forall x_1, x_2 \in C_{i0}$. Fix k and define the set

$$A(k, x_1) := \{y \in W_j : |y \oplus x_1| = k\} \quad \forall x_1 \in C_{i0},$$

then the coefficient of the k th power in the equation is $|A(k, x_1)|$. Furthermore, we have $|A(k, x_1)| = |A(k, x_2)|$ for all $x_1, x_2 \in C_{i0}$.

Notice that C_{j0} was created with the property that $\sum_{x_1 \in C_{i0}} s^{|x_1 \oplus y_1|} = \sum_{x_1 \in C_{i0}} s^{|x_1 \oplus y_2|}$ for all $y_1, y_2 \in C_{j0}$. Define the set

$$B(k, x_1) := \{y_1 \in C_{j0} : |x_1 \oplus y_1| = k\} \quad \forall x_1 \in C_{i0},$$

and its mirror set

$$D(k, y_1) := \{x_1 \in C_{i0} : |x_1 \oplus y_1| = k\} \quad \forall y_1 \in C_{j0}.$$

Notice that $A(k, x_1) \cap C_{j0} = B(k, x_1)$ for all $x_1 \in C_{i0}$.

Now we will use the property $|D(k, y_1)| = |D(k, y_2)|$ for all $y_1, y_2 \in C_{j0}$ to prove that $|B(k, x_1)| = |B(k, x_2)|$ for all $x_1, x_2 \in C_{i0}$. Let $\phi : C_{j0} \rightarrow C_{j0}$ be a bijective map on C_{j0} such that $\phi(y_1) = y_2$. Pick a bijective map $\pi : C_{i0} \rightarrow C_{i0}$ that maps elements of $D(k, y_1)$ to elements of $D(k, \phi(y_1))$. Now, we will show that $y_1 \in B(k, x_1)$ if and only if $\phi(y_1) \in B(k, \pi(x_1))$. Now $y_1 \in B(k, x_1) = A(k, x_1) \cap C_{j0}$, so this is equivalent to $x_1 \in D(k, y_1)$, which in turn is true if and only if $\pi(x_1) \in D(k, \phi(y_1))$, or if and only if $\phi(y_1) \in A(k, \pi(x_1))$. Then since $\phi(y_1) \in C_{j0}$, we have that $\phi(y_1) \in B(k, \pi(x_1))$.

This proves that $|B(k, x_1)| = |B(k, x_2)|$ for all $x_1, x_2 \in C_{i0}$. Therefore we have

$$\sum_{y \in C_{j0}} s^{|y \oplus x_1|} = \sum_k |B(k, x_1)| s^k \quad \forall x_1 \in C_{i0}.$$

Therefore, we have the invariant that

$$\sum_{y \in C_{j0}} s^{|y \oplus x_1|} = \sum_{y \in C_{j0}} s^{|y \oplus x_2|} \quad \forall x_1, x_2 \in C_{i0}$$

□

Theorem 7. (*Uniqueness of Solution.*) *The Maximum Ensemble Algorithm finds the unique solution to the Maximum Ensemble Problem.*

Proof. The partitioning algorithm produces a partition that respects the emission partition, since it begins with the partition given by the emission partition and sub-partitions it. The algorithm also produces partitions that respect the Markov property, since it iteratively sub-partitions the emission partition until the Markov property is satisfied. Notice that the algorithm is guaranteed to find such a partition since the trivial partition, i.e. the original state space, satisfies the Markov property. Since partition sets are only divided if they violate the Markov property, the algorithm necessarily finds an optimal partition. Only the proof of uniqueness remains.

By Lemma 5 the solution is invariant to the order in which the bipartitions are made, since any x_1, x_2 which violate the Markov property must be put into separate sets of the partition at some point. Indeed, by Lemma 6 we know that once C_{i0} is created, it is never partitioned again. Since we begin with a unique partition, the emission partition, the sequence of C_{i0} , created by different calls to Algorithm 1, will be the final sets in the partition, up to reordering. Therefore the Maximum Ensemble Algorithm finds the unique partition which is the solution to the Maximum Ensemble Problem. □

4 Equivalence

Now, using the uniqueness of a partition as the solution to the Maximum Ensemble Problem, we can prove equivalence of the Maximum Ensemble Problem and the Maximum Isometry Group Problem.

Theorem 8. (*Equivalence of Maximum Ensemble Problem and Maximum Isometry Group Problem*) *A partition $\{W_1, W_2, \dots, W_k\}$ is a solution to the Maximum Ensemble Problem if and only if there is an isometry group G that is a solution to the Maximum Group Problem having orbits $\Omega(G)$ equivalent to the partition: for all ω , we have $\omega \in \Omega(G)$ if and only if there exists a set in the partition W_j such that $W_j = \omega$.*

Proof. First, due to Corollary 3, we know that only isometry groups satisfy the Markov property. Any partition which is a solution for the Maximum Ensemble Problem is also, in particular, the orbits of a group of isometries, G . Assume that G is not the maximal isometry group. Because, if not, then there must be some isometry which can be added. And, if it were added, it would join two orbits into one. Therefore joining two sets of the partition into one, which contradicts the assumption that the partition was maximal. Furthermore, since G satisfies the emission property, its orbits must be a subpartition of the emission partition. There is no other group G' with larger size, since the solution to the Maximum Ensemble Problem is unique (Theorem 7). The Maximum Ensemble Problem is a solution to the Maximum Group Problem.

For the converse we argue by contrapositive. Assume that partition $\{W_1, \dots, W_k\}$ is not a solution to the Maximum Ensemble Problem but that it satisfies Equation 3. Then there must exist some partition $\{V_1, \dots, V_l\}$ such that $\sum_{i=1}^k |W_i| < \sum_{j=1}^l |V_j|$. This inequality is strict by the uniqueness Theorem 7. There must exist some i, i' , and j , such that $W_i \subset V_j$ and $W_{i'} \subset V_j$.

By Corollary 3, there are groups G^W and G^V with orbits $\{W_1, \dots, W_k\}$ and $\{V_1, \dots, V_l\}$, respectively. Choose $x_1 \in W_i \cap V_j$ and $x_2 \in W_{i'} \cap V_j$. Then π_{x_1, x_2} from Theorem 1 will be in G^V and not in G^W . Therefore, G^W is not a solution to the Maximal Isometry Group Problem; proving the claim. \square

5 Bootstrapping with Known Isometries

As noted by Geiger et al. [8], there are two types of isometries that can be detected easily. There are the founder isometries and the chain isometries where there is an outbred lineage consisting of multiple ungenotyped generations.

The founder isometries apply only to ungenotyped founders and are switches on the bits for the edges adjacent to the founder. Specifically, if i_1, \dots, i_c are the meiosis bits between the ungenotyped founder and each of the c children of the founder, then the switch is given by the bit vector $X_i = 1$ if $i = i_j$ for some j and $X_i = 0$ otherwise. Since the founder alleles are indistinguishable (due to the missing genotype), we can fix one bit adjacent to the founder and enumerate the other bits adjacent that founder. These founder isometries can be found in $O(n)$ time.

The chain isometries apply to a lineage of l individuals, from oldest to youngest i_1, i_2, \dots, i_l where each individual has exactly one parent from the lineage, one founder parent, one child, and no siblings, except i_l which may have any number of siblings. All individuals except the most recent must be ungenotyped. The isometry is then the permutation on every bit, except the oldest, i.e. $\pi = (1_1 i_2 i_3 \dots i_l)$ Please see Geiger, et al. [8] and Browning and Browning [2] for examples. These chain isometries can be found in $O(n^2)$ time.

It would seem that there are other classes of isometries which can be found quickly, such as the permutations shown in the example in Section 2.3. The exact algorithms for finding other classes of isometries remain an open problem. Furthermore, it is unknown whether all the isometries in the maximal group can be found efficiently.

5.1 Representatives

Let A be a generating set of isometries that generate group $G = \langle A \rangle$, such as the founder and chain isometries. In order to compute the bootstrap maximum ensemble states, We need to obtain the orbits of G acting on \mathcal{H}_n . We can obtain them in $O(k|A|o)$ time where k is the number of orbits and $o = \max_{x \in \mathcal{H}_n} |\omega(x)|$, provided that orbit membership can be checked in constant time.

Let $M = \mathcal{H}_n$ initially. We take any vector x out of M and find its orbit O . Initially let $O = \{x\}$. Now, for every $x \in O$ and every $a \in A$, compute $y = a(x)$. If $y \notin O$, add y to O and remove y from M . Repeat until M is empty.

Following this procedure, we have all of the orbits of G acting on \mathcal{H}_n . For each orbit, we will fix a representative to use in the bootstrap maximal ensemble algorithm.

5.2 Bootstrap Maximal Ensemble

Now that we have k representatives, one from each orbit of group $G = \langle A \rangle$, we can introduce a bootstrap version of the Maximal Ensemble algorithm. In this case, we can compute Equation (3) once per representative.

First, we need to partition our representatives according to the set of the emission partition that they belong to. Consider the emission partition, $\{E_x | \forall x\}$, and partition the representatives into these sets. Also partition \mathcal{H}_n according to the emission partition. These two equivalent partitions define our initial partitions.

Now, we can recursively sub-divide the representatives whenever Equation (3) is violated. Notice that we can compute this equation with x being the representative and ω_j is some set of the current partition of \mathcal{H}_n . Each time we subdivide the partition of the representatives, we need to also subdivide the partition of \mathcal{H}_n in the equivalent fashion. Suppose that we have representative x that we have put into a new set in the representative partition. We obtain the equivalent partition of \mathcal{H}_n by creating a new set containing x and all the vectors $y \in \omega(x)$ the orbit of x under the action of G . The recursive subdivision continues until the Markov property is satisfied.

Since the recursive sub-partitioning at minimum splits sets in two, the number of iterations required is $O(n)$. Checking the Markov properties for each iteration requires $O(k2^n)$ time where k is the number of representatives, since we have to check the $k \times 2^n$ matrix of distances, or sums of distances, between partition elements. So, the total running time is $O(nk2^n)$.

6 Running Times

Notice that the naive calculation of Equation (1) requires $O(k2^n)$ time where $k \leq 2^n$ is the number of sets in the partition and n is the number of meioses in the pedigree. The calculation is as follows, for each set W_i in the partition, choose a representative $x \in W_i$. For each of the sets in the partitions W_j , compute the transition probability $Pr[X_{t+1} \in W_j | X_t = x]$. This last step seems to require enumeration of the inheritance paths.

The running time of the state-space reduction is the running time of the ensemble algorithm and the running time of the transition calculation. It is interesting to note that calculating the transition probabilities in Equation 1 is faster than the HMM forward-backward algorithm having running time $O(m2^{2n})$. This means there is potential to improve the state-space reduction running time, if there is a more efficient maximal ensemble algorithm.

Regardless of whether the over all running time of the state-space reduction is determined by calculating the transition function or the ensemble states, all the algorithms here produces savings when the forward-backward algorithm is run. This is because a k -set partition of the states results in the forward-backward algorithm having $O(mk)$ running time where m is the number of sites. Furthermore, since the original state space has an $O(m2^{2n})$ forward-backward algorithm and the ensemble algorithm is $O(n2^{2n})$, the ensemble algorithm is more efficient when $n < m$ which is typically the case. The bootstrap algorithm is even more efficient having a running time of $O(nk2^n)$.

7 Simulation Results

We simulated pedigrees under a Wright-Fisher model with monogamy where each pair of monogamous individuals has Poisson distributed number of offspring. There are n individuals per generation and λ is the mean number of offspring per monogamous pair. The individuals of interest, I , are the extant individuals, i.e. those in the most recent generation or, equivalently, the nodes with out-degree zero. These pedigrees have no inter-generational mating due to how the Wright-Fisher model is defined. To get a half-sibling pedigree, each edge of the pedigree had 50% chance of have a new parent drawn at random. Since monogamy was not preserved during this random process, the resulting pedigree had half-siblings.

Running the simulation process and the maximal ensemble algorithm 100 times produced Figure 7. The maximal ensemble algorithm produces exponential reductions in the size of the state-space. Whether the relationships have half-siblings seems not to influence the practical applicability of the maximal ensemble algorithm (data not shown).

In practice, the maximal ensemble algorithm seems limited to pedigrees of roughly 14 meioses while the bootstrap maximal ensemble algorithm seems limited to about 18 meioses. Of course, both methods yield the same reduced state space. Given the practical success of the bootstrap maximal ensemble algorithm, we recommend that the bootstrap maximal ensemble algorithm be employed for state-space reduction.

8 Discussion

Even though past efforts at state-space reduction have focused on finding groups of isometries on the edges of the pedigree graph, it is clear that this is an equivalent problem to finding the optimal partition of the state space that respects the Markov property. Although the paper mostly discusses the pedigree state-space, the maximum ensemble algorithm is general to any HMM.

Even if some isometries can be obtained efficiently, for example the founder and chain isometries, computation of the transition probabilities according to Equation 1 seems to require enumeration of the inheritance vectors. The naive algorithm requires $O(k2^n)$ where k is the number of orbits and n is the number of meioses in the pedigree. Due to this fact, and the fact that the forward-backward algorithm for pedigree HMMs has running time $O(m2^{2n})$, it is an advantage to use exponential

transition probabilities of Y_t would allow for a larger reduction in the state-space. In addition, we would hope that the bound on the inequality would guarantee that the deviation of Y_t 's stationary distribution is bounded relative to the stationary distribution of X_t .

Acknowledgments. Many thanks go to Yun Song for suggesting the problem and to Eran Halperin for the random pedigree simulation.

References

- [1] GR Abecasis, SS Cherny, WO Cookson, et al. Merlin-rapid analysis of dense genetic maps using sparse gene flow trees. *Nature Genetics*, 30:97–101, 2002.
- [2] S. Browning and B.L. Browning. On reducing the statespace of hidden Markov models for the identity by descent process. *Theoretical Population Biology*, 62(1):1–8, 2002.
- [3] C. J. Burke and M. Rosenblatt. A Markovian function of a Markov chain. *Ann. Math. Stat.*, 29:1112–1122, 1958.
- [4] G. Coop, X. Wen, C. Ober, et al. High-Resolution Mapping of Crossovers Reveals Extensive Variation in Fine-Scale Recombination Patterns Among Humans. *Science*, 319(5868):1395–1398, 2008.
- [5] K. P. Donnelly. The probability that related individuals share some section of genome identical by descent. *Theoretical Population Biology*, 23(1):34 – 63, 1983.
- [6] John Durbin. *Modern Algebra: An Introduction*. John Wiley and Sons, Inc., 4th edition, 2000.
- [7] M. Fishelson, N. Dovgolevsky, and D. Geiger. Maximum likelihood haplotyping for general pedigrees. *Human Heredity*, 59:41–60, 2005.
- [8] D. Geiger, C. Meek, and Y. Wexler. Speeding up HMM algorithms for genetic linkage analysis via chain reductions of the state space. *Bioinformatics*, 25(12):i196, 2009.
- [9] A. Jacquard. Genetic information given by a relative. *Biometrics*, 28(4):1101–1114, 1972.
- [10] Tommi A. Junttila. New orbit algorithms for data symmetries. *Application of Concurrency to System Design, International Conference on*, 0:175, 2004.
- [11] G. Karigl. A mathematical approach to multiple genetic relationships. *Theoretical Population Biology*, 21:379–393, 1982.
- [12] S. L. Lauritzen and N. A. Sheehan. Graphical models for genetic analysis. *Statistical Science*, 18(4):489–514, 2003.
- [13] Louise Lorentsen and Lars Michael Kristensen. Exploiting stabilizers and parallelism in state space generation with the symmetry method. *Application of Concurrency to System Design, International Conference on*, 0:211, 2001.
- [14] M.S. McPeck. Inference on pedigree structure from genome screen data. *Statistica Sinica*, 12(1):311–336, 2002.

- [15] N. Risch and K. Merikangas. The Future of Genetic Studies of Complex Human Diseases. *Science*, 273(5281):1516–1517, 1996.
- [16] E. Sobel and K. Lange. Descent graphs in pedigree analysis: Applications to haplotyping, location scores, and marker-sharing statistics. *American Journal of Human Genetics*, 58(6):1323–1337, 1996.
- [17] E. A. Thompson. Gene identities and multiple relationships. *Biometrics*, 30(4):pp. 667–680, 1974.
- [18] T. Thornton and M.S. McPeck. Case-control association testing with related individuals: A more powerful quasi-likelihood score test. *American Journal of Human Genetics*, 81:321–337, 2007.